WORLD COMET RESEARCH GROUP

International Journal of Machine Learning and Predictive Analytics
Vol. 1, No. 2, 2018
Doi: 10.67228/3142788X/IJMLPA-2018PII3V7R
PP. 01-11

Original Article

Cognitive Digital Twins with Continual Learning for Industry 4.0 Systems

Dr. Harish Patel[1], Dr. Geetha Ramasamy[2]

[1]*Professor, Gujarat University, India.*

[2]*Associate Professor, Madurai Kamaraj University, India.*

Received: 09-03-2018 Revised: 09-19-2018 Accepted: 09-25-2018 Published: 10-05-2018

ABSTRACT

The evolution of Industry 4.0 has driven the adoption of intelligent and autonomous systems capable of adapting to dynamic environments. Among these, Digital Twins (DTs) have emerged as virtual replicas of physical systems that enable real-time monitoring, simulation, and optimization. However, traditional DTs often suffer from static behavior and lack adaptive intelligence, limiting their utility in non-stationary industrial environments. This paper introduces the concept of Cognitive Digital Twins (CDTs) augmented with Continual Learning (CL) capabilities to overcome these limitations. By integrating cognitive architectures and continual learning mechanisms, CDTs are able to learn from streaming data, retain previously acquired knowledge, and adapt to evolving industrial processes without catastrophic forgetting. The paper presents a conceptual framework for CDTs with CL, explores key enabling technologies, and provides use-case scenarios in predictive maintenance, process optimization, and human-machine collaboration. Challenges and research directions for deploying CDTs in real-world Industry 4.0 ecosystems are also discussed.

KEYWORDS

Cognitive Digital Twin, Continual Learning, Industry 4.0, Adaptive Systems, Catastrophic Forgetting, Cyber-Physical Systems, Smart Manufacturing, Industrial AI, Edge Intelligence, Human-Machine Collaboration.

ISSN: 3142-788X
PP: 01-11.

1. INTRODUCTION

1.1. Background on Industry 4.0 and Digital Twins

Industry 4.0 represents the fourth industrial revolution, characterized by the convergence of cyber-physical systems, the Industrial Internet of Things (IIoT), big data, and artificial intelligence (AI) to create highly interconnected, intelligent, and autonomous manufacturing and production environments. Within this ecosystem, Digital Twins (DTs) have emerged as a foundational technology. A digital twin is a virtual model of a physical asset, system, or process that receives real-time data from sensors to mirror the behavior and status of its real-world counterpart. Initially adopted for predictive maintenance and simulation, DTs have proven valuable in optimizing operations, improving asset utilization, and supporting decision-making. However, as industrial environments grow more dynamic and complex, there is a growing need for digital twins to move beyond static representations and evolve into intelligent systems capable of learning and adaptation.

1.2. Limitations of Traditional Digital Twins

Despite their advantages, traditional digital twins suffer from critical limitations that hinder their applicability in modern industrial contexts. Most conventional DTs rely on predefined models and rules, which makes them inflexible and incapable of dealing with unexpected changes or novel operational conditions. They lack cognitive capabilities, meaning they cannot autonomously interpret new data patterns, update their internal models, or adapt their behavior over time. Their learning processes, if any, are typically offline, requiring manual retraining or human intervention. This static nature limits their ability to function effectively in non-stationary environments where systems degrade, processes evolve, and data distributions shift frequently. Furthermore, traditional DTs are not equipped to retain and generalize knowledge across tasks or contexts, making them unsuitable for applications demanding resilience, continuous improvement, or long-term autonomy.

1.3. The Need for Cognitive and Adaptive Capabilities

To address the evolving requirements of Industry 4.0 systems, there is a pressing need to enhance digital twins with cognitive and adaptive intelligence. Cognitive Digital Twins (CDTs) are envisioned as a next-generation solution, integrating principles of human cognition such as perception, learning, memory, and reasoning into the digital twin paradigm. These enhanced systems can continuously learn from new data, adapt to shifting environmental conditions, and make context-aware decisions autonomously. Such capabilities are critical for real-time fault detection, adaptive process optimization, intelligent control, and effective human-machine collaboration. Importantly, the integration of **continual learning**—where models evolve over time without catastrophic forgetting—enables CDTs to operate robustly in open-ended, non-stationary industrial environments. This transformation is essential for realizing the full potential of autonomous factories, smart logistics, and self-healing systems in Industry 4.0.

1.4. Objectives and Contributions of the Paper

This paper aims to explore the concept of Cognitive Digital Twins enhanced with continual learning mechanisms as a foundational step toward building truly adaptive industrial systems. The

ISSN: 3142-788X
PP: 01-11.

primary objectives are to (1) define and characterize the architecture and components of CDTs, (2) explain how continual learning can be effectively integrated into these systems, (3) demonstrate the applicability of CDTs in various Industry 4.0 scenarios such as predictive maintenance and adaptive control, and (4) identify the key technical challenges and open research questions in deploying CDTs at scale. The contributions of this work include a comprehensive review of related technologies, a proposed modular architecture for CDTs with continual learning, and a discussion on the practical deployment strategies involving edge-cloud computing and cognitive feedback loops. By addressing these topics, the paper seeks to establish a solid conceptual and technical foundation for the next wave of intelligent digital twin systems in industrial environments.

2. RELATED WORK

2.1. Traditional Digital Twin Architectures

Traditional digital twin architectures typically consist of three core layers: the physical layer (i.e., the physical asset or system), the communication layer (which includes data acquisition via sensors and IoT gateways), and the digital layer (which houses the simulation, analytics, or visualization models). These architectures were primarily designed for static use cases such as condition monitoring, offline simulations, or lifecycle management. The underlying models are often physics-based or derived from historical machine learning models trained offline and updated periodically. Despite their effectiveness in certain applications, these architectures do not accommodate dynamic, real-time learning, and they lack built-in mechanisms for knowledge generalization or long-term adaptability. Most implementations treat the digital twin as a passive mirror rather than an active, evolving cognitive agent.

2.2. Cognitive Computing in Industrial Systems

Cognitive computing, which draws from fields like artificial intelligence, neuroscience, and human-computer interaction, aims to simulate human thought processes in complex systems. In industrial contexts, cognitive computing has been explored in applications such as intelligent robotics, autonomous systems, and decision support platforms. These systems typically incorporate machine perception, knowledge representation, learning, and automated reasoning. While cognitive computing has made inroads in areas like natural language interfaces for operator support or intelligent control systems, its integration into the digital twin framework remains limited. The challenge lies in embedding cognitive capabilities into a system originally designed for monitoring and simulation, requiring a shift from reactive architectures to proactive, self-improving systems.

2.3. Continual Learning in AI

Continual learning (CL), also known as lifelong learning, is a subfield of AI that addresses the need for models to learn from a stream of data over time while retaining previously acquired knowledge. This is particularly challenging due to catastrophic forgetting, where new learning overwrites old knowledge. CL techniques are broadly categorized into regularization-based methods (which constrain updates to important weights), replay-based methods (which revisit past data or generate synthetic memories), and parameter isolation techniques (which allocate different model

ISSN: 3142-788X
PP: 01-11.

parts to different tasks). In the context of industrial AI, continual learning is critical for adapting to changing environments, concept drift, and evolving task requirements. However, its adoption in operational digital twin systems is still nascent, largely due to concerns about computational overhead, data privacy, and system stability.

2.4. Existing Integrations of AI with Digital Twins

There has been growing interest in integrating AI techniques with digital twins to enhance their predictive and diagnostic capabilities. Applications range from machine learning-based fault detection and process optimization to reinforcement learning for control systems. These integrations typically involve feeding sensor data into predictive models or optimization algorithms to provide insights or recommendations. However, most AI-enhanced DTs use static models that are retrained offline, with limited real-time learning or adaptability. Few implementations employ continual learning or cognitive feedback mechanisms, and even fewer demonstrate the ability to generalize across tasks or maintain performance in the face of changing data distributions. This indicates a significant gap between current implementations and the vision of truly intelligent, adaptive digital twins.

2.5. Gaps in Current Literature

Despite substantial progress in the individual domains of digital twins, cognitive computing, and continual learning, their integration remains poorly explored. Most existing research treats these areas in isolation, and there is a lack of unified frameworks that combine real-time data integration, adaptive intelligence, and continuous learning within a digital twin system. Furthermore, empirical studies on deploying such systems in real-world industrial settings are scarce, especially those that address scalability, edge-cloud orchestration, or human-machine interaction. There is also limited understanding of the trade-offs involved in applying continual learning techniques to high-frequency industrial data streams. These gaps underscore the need for a holistic approach to designing Cognitive Digital Twins that can operate autonomously, learn continually, and scale effectively across diverse Industry 4.0 environments.

3. COGNITIVE DIGITAL TWINS: DEFINITION AND ARCHITECTURE

3.1. Definition of Cognitive Digital Twins

Cognitive Digital Twins (CDTs) are an advanced evolution of traditional digital twins that not only replicate the physical state and behavior of industrial systems but also possess the ability to reason, learn, and adapt autonomously over time. While a traditional digital twin acts as a passive virtual model, a cognitive digital twin embeds artificial intelligence (AI) to actively interpret data, simulate complex decision processes, and update its internal model based on new experiences. In essence, a CDT goes beyond static simulation and monitoring to provide a continuously learning, decision-supporting digital entity that behaves in a manner akin to human cognition. This shift is essential in Industry 4.0 environments where systems must operate under dynamic conditions, make timely decisions, and collaborate intelligently with humans and machines.

ISSN: 3142-788X
PP: 01-11.

3.2. Core Components

3.2.1. Perception

Perception in CDTs refers to the system's ability to sense, interpret, and contextualize information from its physical counterpart and surrounding environment. This involves a continuous stream of multi-modal data — such as temperature, vibration, pressure, or human input — collected via IoT sensors and edge devices. The perceptual module preprocesses and filters this data, identifying patterns, anomalies, or contextual cues that may signify changing system states or operational risks. Perception in CDTs is not limited to passive data intake but also includes situational awareness, which enables the system to interpret inputs in relation to goals and conditions. By mimicking human-like awareness, perception lays the groundwork for higher-level cognitive functions like learning and decision-making.

3.2.2. Learning and Memory

Learning and memory are central to the cognitive capacity of digital twins, enabling them to evolve their internal models over time and avoid static behavior. The learning module uses machine learning algorithms — including supervised, unsupervised, and continual learning techniques — to derive insights from historical and real-time data. Memory, in this context, refers to both the storage of raw data and the retention of learned patterns or models. The memory system is typically divided into short-term memory, which handles immediate operational knowledge, and long-term memory, which retains generalized knowledge across time. These components allow CDTs to continuously adapt to new situations while maintaining a stable understanding of past behavior, facilitating predictive capabilities and anomaly recognition that improve with time.

3.2.3. Reasoning and Decision-Making

Reasoning and decision-making allow the CDT to infer conclusions and select actions based on current data, learned experiences, and contextual knowledge. This component uses techniques from symbolic AI, probabilistic inference, and reinforcement learning to assess different operational scenarios and simulate potential outcomes. For instance, in a manufacturing setting, a CDT can reason about production bottlenecks, evaluate mitigation strategies, and recommend optimal scheduling changes. The reasoning engine also allows CDTs to explain their decisions, fostering transparency and trust in human-machine interactions. Importantly, this capability transforms CDTs from mere data reflectors into proactive agents capable of strategic interventions and decision support under uncertainty.

3.3. Cognitive Loop Integration in Digital Twins

The cognitive loop is the architectural backbone that connects perception, learning, memory, and reasoning into a cohesive, self-improving system. This loop follows a continuous cycle: the CDT perceives the environment, processes and learns from the incoming data, stores useful knowledge, reasons about operational decisions, and feeds the resulting actions or recommendations back into the physical system or human operator. Each iteration of the loop enhances the CDT's internal models and prepares it to deal more effectively with similar situations in the future. The integration of this cognitive loop ensures that the CDT is not a one-time model but a living digital entity that evolves in synchrony with its physical twin. This integration is particularly vital in volatile industrial settings where continuous improvement and responsiveness are key to operational excellence.

3.4. Comparison with Traditional Digital Twins

Traditional digital twins are largely static or semi-dynamic representations of physical systems, designed primarily for visualization, basic simulation, and historical data analysis. Their intelligence is often limited to rule-based analytics or offline machine learning models that do not adapt in real time. In contrast, cognitive digital twins are dynamic, intelligent, and capable of autonomous evolution. They incorporate AI-driven mechanisms to interpret sensor data in real time, adapt their models as new data arrives, and make context-aware decisions. Unlike traditional DTs, which typically require manual updates and recalibration, CDTs possess a self-learning architecture that allows them to update themselves autonomously. Moreover, CDTs are better equipped to handle concept drift, nonlinear behaviors, and unexpected operational scenarios — all of which are increasingly common in Industry 4.0 environments. This makes CDTs far more suitable for applications that demand resilience, adaptability, and high levels of automation.

4. CONTINUAL LEARNING: FOUNDATIONS AND APPROACHES

4.1. Definition and Importance in Industrial Environments

Continual learning refers to the capability of AI systems to learn incrementally from ongoing streams of data while retaining knowledge acquired from previous experiences. In industrial environments, where processes and operational conditions are highly dynamic and subject to frequent changes due to factors such as equipment wear, changing production requirements, and environmental variations, continual learning becomes a critical feature. It enables Cognitive Digital Twins (CDTs) to update their internal models without the need for complete retraining, ensuring sustained accuracy and adaptability. This adaptive learning is essential for maintaining robust performance in applications like predictive maintenance, anomaly detection, and process optimization, where real-time response to new patterns is crucial.

4.2. Challenges: Catastrophic Forgetting and Data Distribution Shifts

A significant challenge in continual learning is catastrophic forgetting, wherein new learning overwrites or degrades previously learned knowledge, causing a decline in overall system performance. This issue is exacerbated in industrial settings, where long-term historical data is vital for understanding equipment behavior and process trends. Additionally, industrial data often experiences distribution shifts—due to changes in operating conditions, materials, or system configurations—which can mislead models trained under stationary assumptions. Addressing these challenges requires sophisticated learning algorithms capable of balancing plasticity (learning new information) with stability (retaining past knowledge), as well as mechanisms to detect and adapt to evolving data distributions without compromising existing knowledge.

ISSN: 3142-788X
PP: 01-11.

4.3. Techniques

4.3.1. Regularization-Based (e.g., Elastic Weight Consolidation - EWC)

Regularization-based approaches tackle catastrophic forgetting by imposing constraints on model updates, thereby preserving important parameters learned from prior tasks. Elastic Weight Consolidation (EWC) is a prominent technique that uses the Fisher information matrix to quantify parameter importance, penalizing significant deviations during retraining. In the context of CDTs, EWC allows incremental learning with minimal forgetting, making it well-suited for industrial applications where certain operational knowledge must remain stable while adapting to new data.

4.3.2. Replay-Based

Replay-based methods mitigate forgetting by rehearsing previously encountered data or its approximations alongside new data during model updates. This can involve storing representative samples (experience replay) or using generative models to synthesize past data (generative replay). These techniques enable the model to maintain performance across diverse tasks and evolving conditions. For CDTs, replay mechanisms facilitate continual adaptation without discarding historical context, although storage and privacy concerns may necessitate synthetic data generation rather than storing raw industrial data.

4.3.3. Parameter Isolation

Parameter isolation strategies avoid forgetting by allocating dedicated parameters or subnetworks to different tasks, effectively isolating knowledge representations. Techniques like Progressive Neural Networks dynamically expand the model to accommodate new tasks without overwriting previous knowledge. While effective in preventing interference, parameter isolation may lead to increased model complexity and resource demands, posing challenges for deployment in resource-constrained industrial edge devices.

4.4. Suitability of Approaches for CDT Systems

Choosing the right continual learning approach for CDTs depends on the industrial context, data characteristics, and system constraints. Regularization methods like EWC are appealing for their balance between efficiency and retention but may struggle with highly non-stationary data. Replay-based techniques excel in handling distribution shifts but raise concerns about data privacy and storage. Parameter isolation offers robust forgetting prevention but may not scale well with numerous evolving tasks. Hence, hybrid strategies combining multiple approaches or adaptive selection based on task demands are often necessary to build effective, scalable CDT systems that can operate reliably in complex industrial settings.

5. INTEGRATING CONTINUAL LEARNING IN CDTS

5.1. Proposed Architecture for CDTs with CL

The proposed architecture for Cognitive Digital Twins integrated with continual learning encompasses multiple layers that orchestrate data acquisition, processing, learning, and decision-making. At the base is the physical system interfaced through IoT sensors that stream real-time data.

ISSN: 3142-788X
PP: 01-11.

The edge layer handles initial preprocessing and lightweight inference to ensure low-latency responsiveness. A cloud or centralized layer manages deeper model training and knowledge consolidation across distributed twins. Continual learning modules operate across these layers, enabling incremental model updates and knowledge retention. The architecture also integrates a cognitive feedback loop, where the CDT uses reasoning to refine data interpretation and adapt learning parameters dynamically, ensuring the digital twin evolves alongside its physical counterpart.

5.2. Data Flow and Model Update Strategies

Data flow in this architecture follows a continuous loop: raw sensory data is collected and filtered at the edge, relevant features are extracted, and insights are generated through on-device models. Periodically, updated data batches or summaries are transmitted to the cloud for retraining and consolidation of the CDT's global model. Model updates employ continual learning algorithms to incorporate new information without discarding prior knowledge. Incremental updates allow the CDT to respond to immediate changes while maintaining long-term memory. The system dynamically schedules updates based on data novelty, computational resources, and operational priorities to balance accuracy with efficiency.

5.3. Knowledge Retention and Adaptation Mechanisms

To ensure effective knowledge retention, the CDT employs mechanisms such as regularization penalties to protect critical parameters, replay buffers to revisit past experiences, and task-specific memory modules. Adaptation mechanisms monitor changes in input data distributions and trigger learning rate adjustments, model expansions, or selective retraining as needed. These mechanisms collectively maintain a stable yet flexible knowledge base, allowing the CDT to generalize across operational scenarios and avoid catastrophic forgetting while being sensitive to evolving system states.

5.4. Cloud-Edge Deployment Considerations

Deployment of CDTs with continual learning involves a hybrid cloud-edge infrastructure. Edge devices handle time-sensitive data processing and initial model inference to meet real-time constraints. The cloud provides scalable resources for more intensive training, historical data management, and cross-system knowledge sharing. This division addresses latency, bandwidth, and privacy concerns inherent in industrial systems. However, orchestrating model synchronization, ensuring consistency, and managing security across these tiers present significant challenges, requiring robust communication protocols and adaptive learning schedules that respect the heterogeneous and distributed nature of Industry 4.0 environments.

6. USE CASES AND APPLICATIONS

6.1. Predictive Maintenance

In predictive maintenance, CDTs equipped with continual learning can monitor equipment health by analyzing sensor data streams and detecting early signs of wear or failure. Their adaptive

ISSN: 3142-788X
PP: 01-11.

learning capabilities enable them to update failure models as machines age or operating conditions change, improving the timeliness and accuracy of maintenance alerts and reducing downtime and costs.

6.2. Adaptive Process Control

CDTs facilitate adaptive process control by continuously learning from production data to optimize control parameters in real time. They can detect shifts in process dynamics and adjust operational settings to maintain quality and efficiency, thereby reducing waste and energy consumption. Continual learning ensures the control models evolve with new product variants or equipment upgrades.

6.3. Quality Assurance and Anomaly Detection

Quality assurance benefits from CDTs by enabling ongoing learning of normal operational patterns and rapid identification of anomalies. Continual learning allows the system to adapt to new defect types or changing inspection criteria, minimizing false positives and enhancing product reliability.

6.4. Human-in-the-loop Decision Support

CDTs provide augmented decision support by integrating human expertise with AI insights. Continual learning enables the system to incorporate operator feedback, improving model accuracy and interpretability over time. This collaboration enhances trust, reduces cognitive load, and accelerates problem-solving in complex industrial scenarios.

6.5. Case Study (Hypothetical Example)

Consider a smart manufacturing line where a CDT monitors robotic arms and conveyor systems. The CDT continuously learns from sensor inputs to predict joint failures and adjust robot trajectories dynamically. Over months, it adapts to different product lines and varying environmental conditions, significantly reducing unplanned stoppages and optimizing throughput. Operator interventions provide corrective feedback that the CDT incorporates, resulting in progressively improved decision-making and operational resilience.

7. CHALLENGES AND OPEN RESEARCH PROBLEMS

7.1. Data Privacy and Security in CL

Continual learning requires ongoing data collection and model updates, raising concerns about the privacy of sensitive industrial data and protection against cyber-attacks. Ensuring secure data transmission, anonymization, and compliance with industry regulations are critical challenges for deploying CDTs in practice.

7.2. Real-Time Performance Constraints

Industrial systems often demand real-time or near-real-time responses, which can conflict with the computational intensity of continual learning algorithms. Balancing model complexity, inference speed, and update frequency remains an open problem, especially for resource-constrained edge devices.

7.3. Evaluation Metrics for Cognitive Performance

Standardized metrics to evaluate the cognitive capabilities of CDTs—including adaptability, retention, reasoning accuracy, and robustness—are lacking. Developing meaningful benchmarks and validation protocols is essential for comparing approaches and ensuring reliable deployment.

7.4. Standardization and Interoperability

The heterogeneous nature of Industry 4.0 environments necessitates interoperability among diverse devices, platforms, and protocols. Standard frameworks for cognitive digital twins and continual learning integration are underdeveloped, impeding widespread adoption and scalability.

7.5. Trust, Transparency, and Explainability

For CDTs to be accepted in safety-critical industrial contexts, their decisions must be transparent and explainable. Achieving human-understandable reasoning in complex continual learning systems is a significant challenge that requires advances in interpretable AI and human-centered design.

8. CONCLUSION

This paper has presented a comprehensive exploration of Cognitive Digital Twins enhanced with continual learning for Industry 4.0 systems. By defining their architecture, elucidating continual learning foundations, and detailing integration strategies, the work highlights the transformative potential of CDTs to enable adaptive, intelligent industrial operations. Through illustrative use cases and a candid discussion of open challenges, the paper sets the stage for future research that bridges AI, digital twin technology, and industrial automation. The adoption of such systems promises increased efficiency, resilience, and autonomy in manufacturing, signaling a significant leap forward in realizing the vision of Industry 4.0.

REFERENCES

[1] Grieves, M., & Vickers, J. (2017). Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. *Transdisciplinary Perspectives on Complex Systems*, 85-113.

[2] Li, Z., & Hoiem, D. (2018). Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12), 2935-2947.

[3] Kirkpatrick, J., Pascanu, R., Rabinowitz, N., et al. (2017). Overcoming catastrophic forgetting in neural networks. *PNAS*, 114(13), 3521-3526.

[4] Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4), 44.

[5] Susto, G. A., Schirru, A., Pampuri, S., McLoone, S., & Beghi, A. (2015). Machine learning for predictive maintenance: A multiple classifier approach. *IEEE Transactions on Industrial Informatics*, 11(3), 812-820.

[6] Sundararajan, M., & Grandvalet, Y. (2017). Continual learning in robotics: A review. *Robotics and Autonomous Systems*, 100, 91-102.

[7] Grieves, M. (2014). Digital Twin: Manufacturing Excellence through Virtual Factory Replication. Florida Institute of Technology.

ISSN: 3142-788X
PP: 01-11.

[8] Lee, J., Kao, H. A., & Yang, S. (2014). Service innovation and smart analytics for Industry 4.0 and big data environment. Procedia CIRP, 16, 3–8.
[9] Monostori, L. (2014). Cyber-physical production systems: Roots, expectations and R&D challenges. Procedia CIRP, 17, 9–13.
[10] Lee, J., Bagheri, B., & Kao, H. A. (2015). A cyber-physical systems architecture for Industry 4.0-based manufacturing systems. Manufacturing Letters, 3, 18–23.
[11] Kagermann, H., Wahlster, W., & Helbig, J. (2013). Recommendations for Implementing the Strategic Initiative INDUSTRIE 4.0. Acatech.
[12] Hermann, M., Pentek, T., & Otto, B. (2016). Design principles for Industrie 4.0 scenarios. In Proceedings of the 49th Hawaii International Conference on System Sciences (HICSS) (pp. 3928–3937).
[13] Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep Learning. MIT Press.
[14] Russell, S., & Norvig, P. (2016). Artificial Intelligence: A Modern Approach (3rd ed.). Pearson.
[15] Canedo, A. (2016). Industrial IoT lifecycle via digital twins. In Proceedings of the International Conference on Hardware/Software Codesign and System Synthesis.